

The One-Way Communication Complexity of Group Membership

Scott Aaronson
 aaronson@csail.mit.edu
 Massachusetts Institute of Technology

Alexander Russell
 acr@cse.uconn.edu
 University of Connecticut

François Le Gall
 legall@qci.jst.go.jp
 Japan Science and Technology Agency

Seiichiro Tani
 tani@theory.br1.ntt.co.jp
 Japan Science and Technology Agency
 and NTT Corporation

February 21, 2009

Abstract

This paper studies the one-way communication complexity of the *subgroup membership problem*, a classical problem closely related to basic questions in quantum computing. Here Alice receives, as input, a subgroup H of a finite group G ; Bob receives an element $x \in G$. Alice is permitted to send a single message to Bob, after which he must decide if his input x is an element of H . We prove the following upper bounds on the classical communication complexity of this problem in the bounded-error setting:

1. The problem can be solved with $O(\log |G|)$ communication, provided the subgroup H is normal.
2. The problem can be solved with $O(d_{\max} \cdot \log |G|)$ communication, where d_{\max} is the maximum of the dimensions of the irreducible complex representations of G .
3. For any prime p not dividing $|G|$, the problem can be solved with $O(d_{\max} \cdot \log p)$ communication, where d_{\max} is the maximum of the dimensions of the irreducible \mathbb{F}_p -representations of G .

1 Introduction

Background The power of quantum computing in various settings has been gradually clarified by many researchers: some problems can be solved on quantum computers much more efficiently than on classical computers, while others cannot. One computational model that has been extensively studied in this light is the communication complexity model. In particular, *one-way communication* is one of the simplest settings but it has rich connections to areas such as information theory, coding theory, on-line computing, and learning theory. Therefore, its quantum version has then been the target of intensive research [Aar05, INRY07, Kla07, GKK⁺07].

Let $f : X \times Y \rightarrow \{0, 1\}$ be a Boolean function, where X and Y are arbitrary sets. The one-way communication task associated to f is the following: Alice has an input $x \in X$, Bob has an input $y \in Y$ and the goal is for Bob to output $f(x, y)$. The assumption here is that only one message can be sent, from Alice to Bob, and the communication cost of a protocol is the number of bits of this message on the worst-case input. We say that a protocol for f has completeness error ε if it outputs 1 with probability at least $1 - \varepsilon$ whenever $f(x, y) = 1$, and soundness error δ if it outputs 0 with probability at least $1 - \delta$ whenever $f(x, y) = 0$. The

one-way classical bounded-error communication complexity of f , denoted by $R^1(f)$, is the minimum, over all protocols P for f with completeness and soundness error $1/3$, of the communication cost of P . The one-way quantum bounded-error communication complexity of f , denoted by $Q^1(f)$, is defined similarly, but a quantum message can be used this time from Alice to Bob, and the number of qubits of the message is considered (in this paper we suppose that there is no prior entanglement and no shared randomness between Alice and Bob). Obviously for any function f , the relation $Q^1(f) \leq R^1(f) \leq \lceil \log_2 |X| \rceil$ holds.

One of the main open problems in quantum communication complexity is to understand how large the gap between R^1 and Q^1 can be. For partial functions (functions restricted to some domain $R \subset X \times Y$ or, equivalently, functions with a promise on their inputs), an exponential separation between these two quantities has been shown recently in [GKK⁺07]. However the situation for total functions is far less clear: the largest gap known is an asymptotic factor of 2 [Win04].

In the exact setting, i.e., the setting where no error and no giving up are allowed, the quantum and classical one-way communication complexities are known to be the same for any total function [Kla07]. In the unbounded-error setting, i.e., the setting where any error probability less than $1/2$ is allowed, it is known that the gap is exactly a factor 2 for both partial and total functions [INRY07]. Although bounded-error is a notion between the exact and unbounded-error, we stress that the bounded-error setting usually behaves quite differently from the other two in the case of total functions, e.g., for two-way communication there is a quadratic gap in the bounded-error setting [KS92, AA05] whereas in the exact setting no gap is known and, in the unbounded-error setting, the gap is again exactly a factor 2 [INRY07].

Note also that for total functions in the bounded-error setting, quadratic gaps are known in the two-way model [KS92, AA05] and exponential gaps are known in the simultaneous message-passing model [NS96, BCWdW01]; and these models are respectively stronger and weaker than the one-way model. Thus, whether a superlinear gap between R^1 and Q^1 can be achieved for some total function is an intriguing question.

The subgroup membership function Many of the problems for which quantum computation is more powerful than classical computation have group-theoretic structure. In particular, Watrous [Wat00] has used the subgroup membership problem (as a computational problem) to separate the complexity classes MA and QMA relative to an oracle. Inspired by Watrous's work [Wat00], we propose the subgroup membership function as a candidate to show a superlinear gap between R^1 and Q^1 . Let G be any finite group, and let \mathcal{H}_G denote the set of subgroups of G . Then the subgroup membership function for G , denoted by MEMB_G , is the function with domain $\mathcal{H}_G \times G$ such that

$$\text{MEMB}_G(H, y) = \begin{cases} 1 & \text{if } y \in H \\ 0 & \text{if } y \notin H. \end{cases}$$

For any group G , the upper bound $|\mathcal{H}_G| \leq 2^{(\log_2 |G|)^2}$ follows easily from the fact that any subgroup of G is generated by at most $\log_2 |G|$ elements.¹ Furthermore, there exist families of groups G such that $|\mathcal{H}_G| = 2^{\Omega((\log |G|)^2)}$: for example, the abelian groups $G = \mathbb{Z}_2^r$ with $r \geq 1$. Thus there exist groups G for which the “trivial protocol,” wherein Alice simply sends Bob the name of her subgroup, requires $\Theta((\log |G|)^2)$ communication. The one-way classical communication complexity of the function MEMB_G was previously considered by Miltersen et al. [MNSW98], who showed that for the family of groups $G = \mathbb{Z}_2^r$, any one-way protocol with perfect soundness and completeness error $1/2$ requires $\Omega((\log |G|)^2)$ -bit communication. For certain groups G , we conjecture that $\Omega((\log |G|)^2)$ -bit communication is needed even if the completeness and soundness errors are both $1/3$.

¹Borovik, Pyber, and Shalev [BPS96] have improved this naive bound to $|G|^{(1/4+o(1))\log_2 |G|}$.

On the other hand, there is a simple *quantum* one-way protocol, using $O(\log |G|)$ -bit communication, by which Bob can compute MEMB_G with perfect completeness and constant soundness for any group G . In this protocol—inspired by [Wat00]—Alice sends the quantum state $|H\rangle = |H|^{-1/2} \sum_{h \in H} |h\rangle$. Bob then creates the state $\frac{1}{\sqrt{2}}(|H\rangle|0\rangle + |yH\rangle|1\rangle)$ where $|yH\rangle = |H|^{-1/2} \sum_{h \in H} |yh\rangle$, applies a Hadamard gate on the last register, and measures it in the basis $\{|0\rangle, |1\rangle\}$ to decide which of $|H\rangle = |yH\rangle$ and $\langle H|yH\rangle = 0$ holds.

Thus, proving that there exists a family of groups G such that $R^1(\text{MEMB}_G) = \Omega((\log |G|)^2)$ would lead to a quadratic separation between R^1 and Q^1 for a total function. In other words, a major objective has been to prove a 2-sided-error version of the lower bound by Miltersen et al. [MNSW98] mentioned above. Apart from the goal of proving a separation between R^1 and Q^1 , we believe that the communication complexity of deciding subgroup membership is interesting in itself, since the latter is a key task in most group-theoretic computational problems.

Overview of our results In this paper we present three upper bounds on the one-way classical communication complexity of the subgroup membership function:

- We give a classical protocol using $\lceil \log_2 |G| \rceil$ -bit communication, with perfect completeness and constant soundness, for the subgroup membership function in the case where Alice’s subgroup H is normal. This suggests that in order to obtain a separation between R^1 and Q^1 using the subgroup membership problem, one must consider groups with many nonnormal subgroups. We also present a lower bound which is tight for some families of groups. Notice that this situation appears to be similar to the status of the Hidden Subgroup Problem: there exists an efficient quantum algorithm solving the problem in the case where the hidden subgroup is normal [HRTS03]; without the normality condition, however, very little is known. Our results rely on the theory of characters of finite groups and especially on the connection between kernels of irreducible characters and normal subgroups, as did the algorithms of [HRTS03].
- Let p be a prime not dividing $|G|$. Then we show that $R^1(\text{MEMB}_G) = O(d_{\max}^p \cdot \log p)$, where d_{\max}^p is the maximum dimension of an irreducible \mathbb{F}_p -representation of G . This result uses a beautiful characterization of the subspaces of the group algebra $\mathbb{F}_p[G]$ stabilized by H . We remark that for any group G of exponent m (which is to say that $g^m = 1$ for all $g \in G$), we have $d_{\max}^p \leq d_{\max}^0 \text{ord}_m(p)$, where d_{\max}^0 is the maximum dimension of a complex irreducible representation of G and $\text{ord}_m(p)$ is the order of p in \mathbb{Z}_m^* , the multiplicative group of the integers relatively prime to m . In particular, as there is always a prime p of size $O(\log |G|)$ relatively prime to $|G|$, this protocol has complexity no more than $O(d_{\max}^0 \cdot m \cdot \log \log |G|)$.
- Finally, we show that $R^1(\text{MEMB}_G) = O(d_{\max}^0 \cdot \log |G|)$, where d_{\max}^0 is the maximum dimension of an irreducible complex representation of G . This upper bound is obtained by a protocol that mirrors the technique utilized in the modular case by suitably discretizing the vector space \mathbb{C}^d and controlling “geometric expansion” around invariant spaces. One corollary is that any family of groups with an abelian subgroup of constant index has a protocol with complexity $O(\log |G|)$. In particular, for groups such as $G = \mathbb{Z}_2 \ltimes \mathbb{Z}_2^n$, the action of \mathbb{Z}_2 on \mathbb{Z}_2^n being to reverse the order of the coordinates, we have $R^1(\text{MEMB}_G) = O(\log |G|)$.

These results suggest a nontrivial connection between the representation theory of the group G and the subgroup membership problem, and provide natural candidates for which a superlinear separation between $R^1(\text{MEMB}_G)$ and $Q^1(\text{MEMB}_G)$ may be obtained: groups with large irreducible representations and many nonnormal subgroups, e.g., the symmetric group.

2 Preliminaries

We assume the reader is familiar with basic concepts of group theory. Here we introduce some notions from representation theory that we will need. In this paper, G always denotes a finite group and 1 denotes its identity element.

Group representations Let \mathbb{F} be a field whose characteristic does not divide the order of G (so the characteristic of \mathbb{F} could be zero). An \mathbb{F} -representation ρ of G is a homomorphism from G to $\mathrm{GL}(V)$, the group of invertible linear transformations over a vector space V (over the field \mathbb{F}). The *dimension* of ρ is the dimension of V . We say that a representation $\rho : G \rightarrow \mathrm{GL}(V_\rho)$ is *irreducible* if the only subspaces of V_ρ simultaneously fixed by the entire family of linear operators $\rho(g)$ are the trivial ones: $\{\mathbf{0}\}$ and V_ρ .

The *group algebra* $\mathbb{F}[G]$ is the \mathbb{F} -algebra of formal sums

$$\sum_{g \in G} \alpha_g \cdot e_g, \quad \alpha_g \in \mathbb{F},$$

with coordinatewise addition and multiplication defined by linearly extending the rule $e_g \cdot e_h = e_{gh}$. Note that $\mathbb{F}[G]$ has dimension $|G|$ as a vector space over \mathbb{F} . The natural action of G on the group algebra defines the *regular representation*: the action of $x \in G$ on a vector $\mathbf{v} = \sum_{g \in G} \alpha_g \cdot e_g$ in $\mathbb{F}[G]$ is denoted by $x\mathbf{v}$ and defined as

$$x\mathbf{v} = \sum_{g \in G} \alpha_g \cdot e_{xg}.$$

Now, if H is a subgroup of G , let

$$I_H = \{\mathbf{v} \in \mathbb{F}[G] \mid h\mathbf{v} = \mathbf{v} \text{ for all } h \in H\}$$

be the subspace of H -invariant vectors of $\mathbb{F}[G]$. It is easy to check that a vector \mathbf{v} lies in I_H if and only if \mathbf{v} is constant on each left coset of H in G . Let $\mathcal{S}_{G:H}$ be the set of right cosets of H in G . The vectors $\mathbf{v}_S = \sum_{g \in S} e_g$ for $S \in \mathcal{S}_{G:H}$ form a basis of I_H and thus

$$\dim I_H = [G : H],$$

where $[G : H] = |\mathcal{S}_{G:H}| = |G|/|H|$ denotes the index of H in G .

A theorem of Maschke's (see, e.g., [CR06, Ser77]) asserts that $\mathbb{F}[G]$ is semi-simple, i.e., $\mathbb{F}[G]$ can be written as the direct sum of a family of irreducible representations. In this case, a theorem of Wedderburn's [Ser77, CR06] asserts that each irreducible representation appears with multiplicity equal to its dimension:

$$\mathbb{F}[G] = \bigoplus_{\rho \in \mathrm{Irr}(G, \mathbb{F})} V_\rho^{\oplus d_\rho},$$

where $\mathrm{Irr}(G, \mathbb{F})$ denotes the set of (representatives of) all the irreducible \mathbb{F} -representations $\rho : G \rightarrow \mathrm{GL}(V_\rho)$ and d_ρ denotes the dimension of ρ . If $I_H(\rho)$ is the subspace of V_ρ pointwise fixed by H , we see that

$$I_H = \bigoplus_{\rho \in \mathrm{Irr}(G, \mathbb{F})} [I_H(\rho)]^{\oplus d_\rho},$$

and conclude that

$$\sum_{\rho \in \mathrm{Irr}(G, \mathbb{F})} d_\rho \dim I_H(\rho) = [G : H]. \tag{1}$$

Complex characters Let \mathbb{F} be the complex field \mathbb{C} . For any \mathbb{C} -representation ρ of G , the character of ρ is the function $\chi: G \rightarrow \mathbb{C}$ such that $\chi(g) = \text{tr}(\rho(g))$ for any $g \in G$, where tr denotes the trace. Characters are conjugacy class functions: the relation $\chi(gg'g^{-1}) = \chi(g')$ holds for any two elements g, g' of G . Moreover, the value $\chi(1)$ is the dimension of the representation ρ . The kernel of χ , denoted by $\ker(\chi)$, is defined as $\ker(\chi) = \{g \in G \mid \chi(g) = \chi(1)\}$. It is easy to see that $\ker(\chi)$ is a subgroup of G .

A character is said to be irreducible if it is the character of an irreducible representation. Denote by $\text{Char}(G)$ the set of irreducible (complex) characters of G . The relation $\sum_{\chi \in \text{Char}(G)} [\chi(1)]^2 = |G|$ is well-known and implies the inequality $|\text{Char}(G)| \leq |G|$. Let H be a normal subgroup of G . Denote

$$\Lambda_H = \{\chi \in \text{Char}(G) \mid H \leq \ker(\chi)\}.$$

Then the relation

$$\sum_{\chi \in \Lambda_H} [\chi(1)]^2 = [G : H] \tag{2}$$

holds (see, e.g., [Isa76]).

3 Normal subgroups

In this section we give an efficient classical protocol computing the subgroup membership function in the special case where Alice's subgroup H is normal. Our protocol is actually more general: we show that one can decide efficiently membership in the normal closure of H , denoted by \overline{H} (the smallest normal subgroup of G containing H).

The protocol testing normal closure membership, denoted by $\text{NORM}(G)$, is as follows.

Protocol $\text{NORM}(G)$

ALICE'S INPUT: a subgroup $H \in \mathcal{H}_G$.

BOB'S INPUT: an element $y \in G$.

BOB'S OUTPUT: $z \in \{0, 1\}$.

- 1 Alice chooses a random element μ of $\Lambda_{\overline{H}}$ with probability $[\mu(1)]^2 |\overline{H}| / |G|$;
- 2 Alice sends the name of μ to Bob;
- 3 Bob outputs 1 if $\mu(y) = \mu(1)$ and outputs 0 otherwise.

Observe that by equation (2), the weights of Step 1 do indeed determine a probability distribution. Notice that $|\Lambda_{\overline{H}}| \leq |G|$ since $\Lambda_{\overline{H}} \subseteq \text{Char}(G)$ and $|\text{Char}(G)| \leq |G|$. Thus Protocol $\text{NORM}(G)$ can be implemented using $\lceil \log_2 |G| \rceil$ bits of communication. We now show the correctness of this protocol.

Proposition 1. *On any input (H, y) , Protocol $\text{NORM}(G)$ outputs 1 with probability 1 if $y \in \overline{H}$, and outputs 0 with probability at least 1/2 if $y \notin \overline{H}$.*

Proof. If $y \in \overline{H}$, then for any element μ in $\Lambda_{\overline{H}}$ the equality $\mu(y) = \mu(1)$ holds. Then Bob always outputs 1. Protocol $\text{NORM}(G)$ has thus perfect completeness.

Now suppose that $y \notin \overline{H}$. Denote $B = \{\chi \in \Lambda_{\overline{H}} \mid \chi(y) = \chi(1)\}$. The error probability of the protocol is

$$\frac{|\overline{H}|}{|G|} \sum_{\chi \in B} [\chi(1)]^2.$$

To conclude our proof, we now prove that

$$\sum_{\chi \in B} [\chi(1)]^2 \leq \frac{|G|}{2|\bar{H}|}.$$

Let K denote the normal closure of the set $\bar{H} \cup \{y\}$ in G . Remember that the normal closure of a set $S \subseteq G$ is the smallest normal subgroup of G including S , and can be defined explicitly as the subgroup of G generated by all the elements gzg^{-1} for $g \in G$ and $z \in S$. Since $y \notin \bar{H}$ the subgroup \bar{H} is a proper subgroup of K . In particular $|K|/|\bar{H}| \geq 2$. We now claim that $B = \Lambda_K$. Then Equation (2) implies that

$$\sum_{\chi \in B} [\chi(1)]^2 = \sum_{\chi \in \Lambda_K} [\chi(1)]^2 = [G : K] \leq \frac{|G|}{2|\bar{H}|}.$$

The proof of the claim follows. First suppose that χ is an element of Λ_K . Then $\chi(y) = \chi(1)$ and thus $\chi \in B$. Now suppose that χ is an element of B . Then $\bar{H} \cup \{y\} \subseteq \ker(\chi)$. From the basic properties of characters, we conclude that $K \subseteq \ker(\chi)$ and thus $\chi \in \Lambda_K$. \square

Given a finite group G , let \mathcal{H}_G^* be the set of normal subgroups of G . Since for a normal subgroup H of G we have $\bar{H} = H$, we conclude that Protocol $\text{NORM}(G)$ solves the restriction of MEMB_G to the domain $\mathcal{H}_G^* \times G$ (notice that this is still a total function).

Theorem 1. *For any finite group G , the restriction of MEMB_G to the domain $\mathcal{H}_G^* \times G$ can be computed with perfect completeness and soundness error $1/2$ by communicating at most $\lceil \log_2 |G| \rceil$ bits.*

We now show a simple lower bound on the communication complexity of MEMB_G . We first recall the definition of the VC-dimension of a set of functions [VC71].

Definition 1. *Let Σ be a set of Boolean functions over a finite domain Y . We say that a set $S \subseteq Y$ is shattered by Σ if for every subset $R \subseteq S$ there exists a function $\sigma_R \subseteq \Sigma$ such that $\forall y \in S, (\sigma_R(y) = 1 \text{ if and only if } y \in R)$. The largest size of set S over all S shattered by Σ is the VC-dimension of Σ , denoted by $VC(\Sigma)$.*

We say that a subset S of a finite group G is an independent subset of G if, for each $g \in S$, element g cannot be written as any product of elements of $S \setminus \{g\}$. We denote by $\gamma(G)$ the maximal size of an independent subset of G . Notice that, for any finite group G , the inequality $\gamma(G) \leq \log_2 |G|$ holds. We now state our lower bound.

Proposition 2. $Q^1(\text{MEMB}_G) = \Omega(\gamma(G))$. In particular, the family of groups $G = \mathbb{Z}_2^r$ for $r \geq 1$ satisfies $Q^1(\text{MEMB}_G) = \Omega(\log |G|)$.

Proof. For each subgroup $H \in \mathcal{H}_G$, define the function $f_H : G \rightarrow \{0, 1\}$ as $f_H(y) = \text{MEMB}_G(H, y)$ for every $y \in G$. Denote $\Sigma = \{f_H \mid H \in \mathcal{H}_G\}$. A result by Klauck [Kla07] shows that $Q^1(\text{MEMB}_G) \geq (1 - h(1/3)) \cdot VC(\Sigma)$, where h is the binary entropy function.

Let $g_1, \dots, g_{\gamma(G)}$ be distinct elements of G such that $S = \{g_1, \dots, g_{\gamma(G)}\}$ is a subset of independent elements of G . The subset $S \subseteq G$ is shattered by Σ since it is easy to show that, for any subset $R \subseteq S$, the function $f_{\langle R \rangle}$ is such that $\forall y \in S, f_{\langle R \rangle}(y) = 1$ if and only if $y \in R$ (here $\langle R \rangle$ denotes the subgroup generated by the elements in R). Then $VC(\Sigma) \geq \gamma(G)$ and $Q^1(\text{MEMB}_G) \geq (1 - h(1/3)) \cdot \gamma(G)$.

The second part of the proposition follows from the observation that each group \mathbb{Z}_2^r is also a vector space of dimension r over the finite field \mathbb{Z}_2 and, thus, $\gamma(\mathbb{Z}_2^r) = r = \log_2(|\mathbb{Z}_2^r|)$. \square

Proposition 2 shows that, for the family of groups $G = \mathbb{Z}_2^r$, Protocol $\text{NORM}(G)$ is optimal up to a constant factor.

4 Algorithms for groups with small modular representations

In this section we present a protocol computing the group membership function for groups with small modular representations. Let \mathbb{F}_q be a finite field with characteristic p not dividing $|G|$. Our protocol, denoted by $\text{MOD-REP}(G, \mathbb{F}_q)$, is the following.

Protocol $\text{MOD-REP}(G, \mathbb{F}_q)$

ALICE'S INPUT: a subgroup $H \in \mathcal{H}_G$.

BOB'S INPUT: an element $y \in G$.

BOB'S OUTPUT: $z \in \{0, 1\}$.

- 1 Alice chooses a representation $\rho : G \rightarrow \text{GL}(V_\rho)$ in $\text{Irr}(G, \mathbb{F}_q)$ with probability $\frac{|H| \cdot d_\rho \cdot \dim I_H(\rho)}{|G|}$;
- 2 Alice chooses a random vector $\mathbf{v} \in I_H(\rho) \subseteq V_\rho$;
- 3 Alice sends the name of ρ and the vector \mathbf{v} to Bob;
- 4 Bob outputs 1 if $\rho(y)\mathbf{v} = \mathbf{v}$ and outputs 0 otherwise.

Observe that by equation (1), the weights of Step 1 do indeed determine a probability distribution.

We now show the correctness of this protocol.

Theorem 2. *Let G be a finite group and \mathbb{F}_q be a finite field of characteristic p not dividing $|G|$. Protocol $\text{MOD-REP}(G, \mathbb{F}_q)$ computes MEMB_G with perfect completeness and constant soundness error. Its communication complexity is at most $\lceil \log_2 |G| \rceil + d_{\max}^q \cdot \lceil \log_2 q \rceil$ bits, where d_{\max}^q is the maximum dimension of an irreducible \mathbb{F}_q -representation of G .*

Proof. Note that the protocol is clearly complete: if $y \in H$, then Bob always accepts.

To establish soundness, let $y \notin H$ and define $K = \langle H, y \rangle$, the smallest subgroup containing both H and y . Remember that $I_K(\rho)$ denotes the subspace of V_ρ pointwise fixed by K . We see that

$$\mathbb{E}_\rho \left[\frac{\dim I_K(\rho)}{\dim I_H(\rho)} \right] = \sum_{\rho} \frac{|H| d_\rho \dim I_K(\rho)}{|G|} = \frac{|H|}{|K|} = \frac{1}{[K : H]} \leq \frac{1}{2},$$

again by equation (1). Observe, then, that $I_K(\rho) \subseteq I_H(\rho)$ and so

$$\mathbb{E}_\rho \left[\frac{\dim I_K(\rho)}{\dim I_H(\rho)} \right] \geq \Pr[I_K(\rho) = I_H(\rho)].$$

Then

$$\Pr[I_K(\rho) \neq I_H(\rho)] = 1 - \Pr[I_K(\rho) = I_H(\rho)] \geq 1/2.$$

When $I_K(\rho) \neq I_H(\rho)$, the vector \mathbf{v} chosen by Alice has probability no more than $1/q$ to be in $I_K(\rho)$. Then $\rho(y)\mathbf{v} \neq \mathbf{v}$ with constant probability in her choices of ρ and \mathbf{v} .

Since $|\text{Irr}(G, \mathbb{F}_q)| \leq |G|$, the communication complexity of the protocol is at most $\lceil \log_2 |G| \rceil + d_{\max}^q \cdot \lceil \log_2 q \rceil$. \square

In light of the complexity guarantee of the protocol above, it is natural to ask how the dimensions of the irreducible representations of a finite group G compare over various fields and, especially, how the modular case compares to the complex case. When the group algebras involved are semi-simple (as they are in this paper due to our insistence that $p \nmid |G|$), there is a tight connection expressed in the following proposition.

Proposition 3. *Let G be a finite group of exponent m and p be any prime not dividing $|G|$. Then the relation $d_{\max}^p \leq d_{\max}^0 \text{ord}_m(p)$ holds, where d_{\max}^0 is the maximum dimension of a complex irreducible representation of G , d_{\max}^p is the maximum dimension of an irreducible \mathbb{F}_p -representation of G , and $\text{ord}_m(p)$ is the order of p in \mathbb{Z}_m^* , the multiplicative group of the integers relatively prime to m .*

Proof. This is a consequence of the “c-d-e triangle” (see [Ser77]). See Appendix A for a brief discussion. \square

As there always exists a prime p of size $O(\log |G|)$ that does not divide $|G|$, we obtain the following corollary.

Corollary 1. $R^1(\text{MEMB}_G) = O(d_{\max}^0 \cdot m \cdot \log \log |G|)$, where m denotes the exponent of G and d_{\max}^0 is the maximum dimension of a complex irreducible representation of G .

5 Algorithms for groups with small \mathbb{C} -representations

We now focus on the case where the dimensions of the irreducible \mathbb{C} -representations of G is under control. The key idea is to discretize the protocol given in the previous section. To achieve this goal we use the concept of an ε -net of a sphere. (As our nets will lie in the vector spaces acted upon by the irreps of G , we define them as subsets of complex Hilbert spaces.)

Definition 2. *Let V be a finite-dimensional complex Hilbert space. An ε -net of V is a finite family of unit-vectors $N \subseteq V$ so that for every unit-length vector $\mathbf{w} \in V$, there is a vector $\mathbf{n} \in N$ so that $|\langle \mathbf{n}, \mathbf{w} \rangle|^2 > 1 - \varepsilon^2$.*

Proposition 4. *For any $\varepsilon > 0$ and for any complex Hilbert space V of dimension d , there exists an ε -net of size at most $(4/\varepsilon)^{2d}$.*

Proof. For any dimension d and distance $\varepsilon > 0$, there is a set of points $A \subset S^{d-1}$ of cardinality no more than $(4/\varepsilon)^d$ with the property that every point of S^{d-1} has distance no more than ε from some point of A (see, e.g., [Mat02, §3.1]). This yields a set with analogous properties of size no more than $(4/\varepsilon)^{2d-1}$ for the complex d -sphere, which has the same metric as the real $2d-1$ sphere. Note that if \mathbf{v} and \mathbf{w} are two unit vectors of V , we may write $\mathbf{v} = \langle \mathbf{v}, \mathbf{w} \rangle \mathbf{w} + \mathbf{r}$ with $\langle \mathbf{r}, \mathbf{w} \rangle = 0$ in which case, $\|\mathbf{r}\| \leq \|\mathbf{v} - \mathbf{w}\|$. The statement of the proposition follows. \square

Our protocol requires the choice of a sufficiently dense ε -net for each irreducible representation in $\text{Irr}(G, \mathbb{C})$. This choice is independent of the inputs of the protocol and so can be done by Alice and Bob without communication. The protocol is as follows.

Protocol COMP-REP(G, ε)

ALICE’S INPUT: a subgroup $H \in \mathcal{H}_G$

BOB’S INPUT: an element $y \in G$

BOB’S OUTPUT: $z \in \{0, 1\}$.

- 1 Alice and Bob agree on an ε -net N_ρ of V_ρ for each $\rho : G \rightarrow \text{GL}(V_\rho)$ in $\text{Irr}(G, \mathbb{C})$;
- 2 Alice chooses a representation $\rho : G \rightarrow \text{GL}(V_\rho)$ in $\text{Irr}(G, \mathbb{C})$ with probability $\frac{|H| \cdot d_\rho \cdot \dim I_H(\rho)}{|G|}$;
- 3 Alice chooses a random (according to Haar measure) unit length vector $\mathbf{v} \in I_H(\rho) \subseteq V_\rho$;
- 4 Alice sends Bob the name of ρ and the closest vector \mathbf{n} in N_ρ to the vector \mathbf{v} ;
- 5 If $|1 - \langle \rho(y)(\mathbf{n}), \mathbf{n} \rangle| \leq 2\varepsilon$, then Bob outputs 1;
Otherwise $|1 - \langle \rho(y)(\mathbf{n}), \mathbf{n} \rangle| > 2\varepsilon$, and Bob outputs 0.

Observe that by equation (1), the weights at Step 2 do indeed determine a probability distribution on $\text{Irr}(G, \mathbb{C})$. Ideally, at Step 3, Alice would communicate \mathbf{v} to Bob: Bob could then check if $\rho(y)(\mathbf{v}) = \mathbf{v}$ and, if so, would figure that $y \in H$. If $\rho(y)(\mathbf{v}) \neq \mathbf{v}$, Bob would be sure that $y \notin H$, since $I_H(\rho)$ is precisely the fixed space of H . The proof below shows that by sending a sufficiently close approximation to \mathbf{v} , Bob can still answer confidently.

The following theorem states the correctness and the communication complexity of this protocol.

Theorem 3. *There exists a choice of ε_G such that Protocol COMP-REP(G, ε_G) computes MEMB_G with perfect completeness and constant soundness error by communicating $O(d_{\max}^0 \cdot \log |G|)$ bits, where d_{\max}^0 denotes the maximum dimension of an irreducible \mathbb{C} -representation of G .*

Proof. As the name of the representation ρ can be encoded using $\lceil \log_2 |G| \rceil$ bits, the communication complexity of the protocol will be dominated by the number of bits necessary to encode the vector \mathbf{n} . We will show that a choice $\varepsilon = \varepsilon_G = \Omega(1/(|G|^2 \text{poly log } |G|))$ suffices to achieve perfect completeness and constant soundness. According to Proposition 4, such an ε -net can be indexed with $O(d_{\rho} \log |G|)$ bits. This gives our upper bound.

We proceed with the analysis of the completeness and soundness of the protocol.

Completeness Observe that if $y \in H$, then the vector \mathbf{v} chosen by Alice in the protocol is fixed by $\rho(y)$. Recall that Alice sends Bob a vector \mathbf{n} for which $|\langle \mathbf{n}, \mathbf{v} \rangle|^2 \geq 1 - \varepsilon^2$; writing

$$\mathbf{n} = \langle \mathbf{n}, \mathbf{v} \rangle \mathbf{v} + \mathbf{r}$$

(where $\langle \mathbf{r}, \mathbf{v} \rangle = 0$) we have

$$1 = \langle \mathbf{n}, \mathbf{n} \rangle = |\langle \mathbf{n}, \mathbf{v} \rangle|^2 + \langle \mathbf{r}, \mathbf{r} \rangle$$

and $\|\mathbf{r}\| \leq \varepsilon$. Considering that

$$\langle \rho(y)\mathbf{n}, \mathbf{n} \rangle = |\langle \mathbf{n}, \mathbf{v} \rangle|^2 + \langle \rho(y)\mathbf{r}, \mathbf{n} \rangle$$

we conclude that

$$|1 - \langle \rho(y)\mathbf{n}, \mathbf{n} \rangle| = |1 - |\langle \mathbf{n}, \mathbf{v} \rangle|^2 - \langle \rho(y)\mathbf{r}, \mathbf{n} \rangle| \leq (1 - |\langle \mathbf{n}, \mathbf{v} \rangle|^2) + |\langle \rho(y)\mathbf{r}, \mathbf{n} \rangle| \leq \varepsilon^2 + |\langle \rho(y)\mathbf{r}, \mathbf{n} \rangle|.$$

Recall that $\rho(y)$ is unitary, so that $\|\rho(y)\mathbf{r}\| = \|\mathbf{r}\|$. Then, by the Cauchy-Schwarz inequality,

$$|1 - \langle \rho(y)\mathbf{n}, \mathbf{n} \rangle| \leq \varepsilon^2 + \|\mathbf{r}\| \leq \varepsilon^2 + \varepsilon.$$

As $\varepsilon < 1$, we have $\varepsilon^2 + \varepsilon \leq 2\varepsilon$ and it follows that the protocol has perfect completeness.

Soundness We wish to show that for sufficiently small ε ($= 1/\text{poly}|G|$), the protocol has constant soundness. Assume that $y \notin H$ and let $K = \langle H, y \rangle$, the smallest subgroup containing H and y . Our goal will be to show that with constant probability $\langle \mathbf{v}, \rho(y)\mathbf{v} \rangle$ is far from 1, in which case the same can be said of \mathbf{n} so long as ε is sufficiently small.

From equation (1),

$$\mathbb{E}_{\rho} \left[\frac{\dim I_K(\rho)}{\dim I_H(\rho)} \right] = \sum_{\rho} \frac{|H| d_{\rho} \dim I_K(\rho)}{|G|} = \frac{|H|}{|K|} = \frac{1}{[K : H]} \leq \frac{1}{2}.$$

Then, with constant probability, the subspace of $I_H(\rho)$ fixed by y has dimension no more than $2/3 \cdot \dim I_H(\rho)$. We may write the vector $\mathbf{v} \in I_H(\rho)$ as $\mathbf{v} = \mathbf{v}_y + \mathbf{v}'$, where $\mathbf{v}_y \in I_K(\rho)$ and $\mathbf{v}' \in [I_K(\rho)]^{\perp}$, the space

perpendicular to $I_K(\rho)$. We then have $\rho(y)\mathbf{v}_y = \mathbf{v}_y$ and $\mathbf{v}_y \in I_K(\rho) \subset I_H(\rho)$. Now, as \mathbf{v} is chosen uniformly on the unit sphere in V_ρ , we have $\mathbb{E}_{\mathbf{v}}[\|\mathbf{v}_y\|^2] = \dim I_K(\rho) / \dim I_H(\rho)$ and the probability

$\Pr_{\rho, \mathbf{v}}[\|\mathbf{v}'\|^2 \geq 1/6]$ is lower bounded by a constant.² We wish to conclude that, conditioned on the event $\|\mathbf{v}'\|^2 \geq 1/6$, the value

$$\frac{\langle \mathbf{v}', \rho(y)\mathbf{v}' \rangle}{\|\mathbf{v}'\|^2}$$

cannot be too close to 1. We will show, in fact, that the real part is appropriately bounded below 1.

Consider the restriction of the representation $\rho : G \rightarrow \mathrm{GL}(V_\rho)$ chosen by Alice to the subgroup K : specifically, we may decompose V_ρ as an orthogonal direct sum of K -invariant subspaces:

$$V_\rho = \bigoplus_i W_{\sigma_i},$$

where each σ_i is in $\mathrm{Irr}(K, \mathbb{C})$ (but copies of the same irrep may appear several times in the direct sum). In this decomposition, \mathbf{v}_y is precisely the projection of \mathbf{v} into the subspace $\bigoplus_{i: \sigma_i=1} W_{\sigma_i}$ corresponding to the copies of the trivial representation; \mathbf{v}' , on the other hand, lies solely in $\bigoplus_{i: \sigma_i \neq 1} W_{\sigma_i}$. As both \mathbf{v} and \mathbf{v}_y lie in $I_H(\rho)$, the difference \mathbf{v}' does as well and the projection of \mathbf{v}' into each W_{σ_i} is H -invariant (that is, lies in $I_H(\sigma_i)$). With this in mind, we shall upper bound

$$\frac{\Re \langle \mathbf{v}', \rho(y)\mathbf{v}' \rangle}{\|\mathbf{v}'\|^2}$$

by controlling

$$\lambda_y \triangleq \max_{\sigma \neq 1} \max_{\mathbf{w} \in I_H(\sigma)} \frac{\Re \langle \mathbf{w}, \sigma(y)\mathbf{w} \rangle}{\|\mathbf{w}\|^2}$$

taken over all nontrivial irreps σ of K and all H -invariant vectors \mathbf{w} in W_σ . In particular, writing $\mathbf{v}' = \sum_{i: \sigma_i \neq 1} \mathbf{v}'_i$ (with each \mathbf{v}'_i lying in W_{σ_i}), we have $\|\mathbf{v}'\|^2 = \sum_{i: \sigma_i \neq 1} \|\mathbf{v}'_i\|^2$ and

$$\Re \langle \mathbf{v}', \rho(y)\mathbf{v}' \rangle = \sum_{i: \sigma_i \neq 1} \Re \langle \mathbf{v}'_i, \rho(y)\mathbf{v}'_i \rangle \leq \sum_{i: \sigma_i \neq 1} \lambda_y \|\mathbf{v}'_i\|^2 = \lambda_y \|\mathbf{v}'\|^2.$$

Observe that if A is a set of generators for H and \mathbf{w} is an H -invariant vector of W_σ ,

$$\langle \mathbf{w}, \sigma(y)\mathbf{w} \rangle = \langle \mathbf{w}, \sigma(y)S_A \mathbf{w} \rangle$$

where $S_A = S_A^\sigma = \frac{1}{|A|} \sum_{a \in A} \sigma(a)$. Then

$$\lambda_y \leq \max_{\sigma \neq 1} \max_{\mathbf{w} \in W_\sigma} \frac{\Re \langle \mathbf{w}, \sigma(y)S_A \mathbf{w} \rangle}{\|\mathbf{w}\|^2}.$$

(Note that the vector \mathbf{w} is not constrained to be H -invariant in this expression.) If we choose A to be a symmetric generating set (so that $a \in A \Leftrightarrow a^{-1} \in A$) then S_A is self-adjoint and $\sigma(y)$ is unitary so that

$$\max_{\sigma \neq 1} \max_{\|\mathbf{w}\|=1} \Re \langle \mathbf{w}, \sigma(y)S_A \mathbf{w} \rangle = \max_{\sigma \neq 1} \max_{\|\mathbf{w}\|=1} \frac{1}{2} \left[\langle \mathbf{w}, \sigma(y)S_A \mathbf{w} \rangle + \langle \mathbf{w}, S_A \sigma(y^{-1}) \mathbf{w} \rangle \right].$$

²Of course, when $\dim I_K(V_\rho) < 2/3 \dim I_H(V_\rho)$, the random variable $\|\mathbf{v}'\|^2$ possesses much stronger concentration around the expected value than this.

As the operator $\sigma(y)S_A + S_A\sigma(y^{-1})$ is Hermitian, we have

$$\max_{\sigma \neq 1} \max_{\|\mathbf{w}\|=1} \Re \langle \mathbf{w}, \sigma(y)S_A \mathbf{w} \rangle \leq \max_{\sigma \neq 1} \left\| \frac{\sigma(y)S_A + S_A\sigma(y^{-1})}{2} \right\|$$

where $\|\cdot\|$ denotes the operator norm.

In order to control this operator norm, observe that the linear operator $(1/2) [\sigma(y)S_A + S_A\sigma(y^{-1})]$ is precisely given by the left action of the group algebra element

$$[\mathbf{A}, \mathbf{y}] \triangleq \frac{1}{2|A|} \left[\sum_{a \in A} e_{ya} + \sum_{a \in A} e_{ay^{-1}} \right] \in \mathbb{C}[K] \quad (3)$$

on the invariant subspace W_σ of $\mathbb{C}[K]$ corresponding to the representation σ . Alternatively, we may consider the Cayley graph on the group K given by the symmetric generating (multi-)set $yA \cup Ay^{-1}$. The (normalized) adjacency matrix of this Cayley graph is identical to the regular representation evaluated at the group algebra element (3) above. As $yA \cup Ay^{-1}$ is a (symmetric) generating set for K , the operator norm of $\sigma([\mathbf{A}, \mathbf{y}])$ is bounded below 1 for each nontrivial σ (see, e.g., [HLW06]). In order to conclude the proof, we require explicit bounds on this spectral gap.

A result of Erdős and Renyi [ER65] asserts that we may select a set of generators A for H of size $O(\log |H|)$ so that the diameter of the resulting Cayley graph (generated by A over H) is $O(\log |H|)$. Considering that the diameter of A (as generators for H) is $O(\log |H|)$, it is easy to see that the set $yA \cup Ay^{-1}$ induces a Cayley graph on K of diameter no more than $O([K : H] \log |H|)$.

Now we may invoke a theorem of Babai [Bab91] asserting that the second eigenvalue of any (undirected) Cayley graph with degree d and diameter Δ is no more than $d - \Omega(1/\Delta^2)$. (If we normalize the adjacency matrix by degree, the second eigenvalue is no more than $1 - \Omega(1/(d\Delta^2))$.) We conclude that

$$\Re \frac{\langle \mathbf{v}', \rho(y) \mathbf{v}' \rangle}{\|\mathbf{v}'\|^2} \leq \lambda_y \leq \max_{\sigma \neq 1} \left\| \frac{\sigma(y)S_A + S_A\sigma(y^{-1})}{2} \right\| \leq 1 - \Omega\left(\frac{1}{[K : H]^2 \log^3 |H|}\right)$$

and, considering that $\|\mathbf{v}'\|^2 \geq 1/6$, that

$$\Re \langle \mathbf{v}, \rho(y) \mathbf{v} \rangle \leq \|\mathbf{v}_y\|^2 + \frac{\Re \langle \mathbf{v}', \rho(y) \mathbf{v}' \rangle}{\|\mathbf{v}'\|^2} \cdot \|\mathbf{v}'\|^2 \leq 1 - \Omega\left(\frac{1}{[K : H]^2 \log^3 |H|}\right).$$

Finally, Alice's \mathbf{n} can be written $\mathbf{n} = \mathbf{v} + \mathbf{r}$ with $\|\mathbf{r}\| \leq \varepsilon$, in which case

$$|\langle \mathbf{n}, \rho(y) \mathbf{n} \rangle| \leq 1 - \Omega\left(\frac{1}{[K : H]^2 \log^3 |H|}\right) + 3\varepsilon \leq 1 - 2\varepsilon,$$

for $\varepsilon^{-1} = \Omega([K : H]^2 \log^3 |H|)$; thus the protocol is sound. \square

In particular, Theorem 3 shows that, over groups for which d_{\max}^0 is constant, the subgroup membership problem can be solved using $O(\log |G|)$ -bit communication. There is a very beautiful characterization of such groups: a family of groups has representations of bounded degree if and only each group of the family has an abelian subgroups of constant index [Glu85]. We thus obtain the following corollary.

Corollary 2. *Let G be a family of groups each possessing an abelian subgroup of constant index. Then $R^1(\text{MEMB}_G) = O(\log |G|)$.*

Acknowledgments

The authors are grateful to Keith Conrad, Iordanis Kerenidis, and Troy Lee for helpful discussions on this subject.

References

- [AA05] Scott Aaronson and Andris Ambainis. Quantum search of spatial regions. *Theory of Computing*, 1:47–79, 2005.
- [Aar05] Scott Aaronson. Limitations of quantum advice and one-way communication. *Theory of Computing*, 1:1–28, 2005.
- [Bab91] László Babai. Local expansion of vertex-transitive graphs and random generation in finite groups. In *Proceedings of the 23rd annual ACM symposium on Symposium of Theory of Computing*, pages 164–174, 1991.
- [BCWdW01] Harry Buhrman, Richard Cleve, John Watrous, and Ronald de Wolf. Quantum fingerprinting. *Physical Review Letters*, 87(16):167902, 2001.
- [BPS96] Alexander Borovik, Laszlo Pyber, and Aner Shalev. Maximal subgroups in finite and profinite groups. *Transactions of the AMS*, 348(9):3734–3761, 1996.
- [CR06] Charles W. Curtis and Irving Reiner. *Representation theory of finite groups and associative algebras*. American Mathematical Society, 2006.
- [ER65] P. Erdős and A. Rényi. Probabilistic methods in group theory. *Journal d'Analyse Mathématique*, 14:127–138, 1965.
- [GKK⁺07] Dmitry Gavinsky, Julia Kempe, Iordanis Kerenidis, Ran Raz, and Ronald de Wolf. Exponential separations for one-way quantum communication complexity, with applications to cryptography. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 516–525, 2007.
- [Glu85] David Gluck. The largest irreducible degree of a finite group. *Canadian Journal of Mathematics*, 37(4):442–451, 1985.
- [HLW06] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.
- [HRTS03] Sean Hallgren, Alexander Russell, and Amnon Ta-Shma. The hidden subgroup problem and quantum computation using group representations. *SIAM Journal on Computing*, 32(4):916–934, 2003.
- [INRY07] Kazuo Iwama, Harumichi Nishimura, Rudy Raymond, and Shigeru Yamashita. Unbounded-error one-way classical and quantum communication complexity. In *Proceedings of the 34th International Colloquium on Automata, Languages and Programming*, pages 110–121, 2007.
- [Isa76] I. Martin Isaacs. *Character theory of finite groups*. Dover, 1976.

- [Kla07] Hartmut Klauck. One-way communication complexity and the Nečiporuk lower bound on formula size. *SIAM Journal on Computing*, 37(2):552–583, 2007.
- [KS92] Bala Kalyanasundaram and Georg Schnitger. The probabilistic communication complexity of set intersection. *SIAM Journal on Discrete Mathematics*, 5(4):545–557, 1992.
- [Mat02] Jiri Matousek. *Lectures on discrete geometry*. Springer, 2002.
- [MNSW98] Peter Bro Miltersen, Noam Nisan, Shmuel Safra, and Avi Wigderson. On data structures and asymmetric communication complexity. *Journal of Computer and System Sciences*, 57(1):37–49, 1998.
- [NS96] Ilan Newman and Mario Szegedy. Public vs. private coin flips in one round communication games (extended abstract). In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing*, pages 561–570, 1996.
- [Ser77] Jean-Pierre Serre. *Linear representations of finite groups*. Springer, 1977.
- [VC71] Vladimir Vapnik and Alexey Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and Applications*, 16:264–280, 1971.
- [Wat00] John Watrous. Succinct quantum proofs for properties of finite groups. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, pages 537–546, 2000.
- [Win04] Andreas Winter. Quantum and classical message protect identification via quantum channels. *Quantum Information and Computation*, 4(6&7):564–578, 2004.

A Remarks on the relationship between \mathbb{C} and \mathbb{F}_p representations

Let G be a finite group of exponent m (so m is the smallest integer for which $g^m = 1$ for all $g \in G$). We outline a technique for reducing \mathbb{C} -representations of G to \mathbb{F}_p -representations in a manner that preserves irreducibility. For a complete account, see [Ser77]. By a difficult theorem of Brauer (see, e.g., [CR06]), one may always realize a \mathbb{C} -irrep over the field $\mathbb{Q}[\zeta_m]$ where ζ_m is a principal m th root of unity. (It is natural to guess that this might be true, as all eigenvalues of a representation of G are m th roots of unity.) Let $\mathbb{Z}[\zeta_m]$ be the ring of algebraic integers in $\mathbb{Q}[\zeta_m]$ (it so happens that in this cyclotomic case $\mathbb{Z}[\zeta_m]$ is indeed the ring of algebraic integers). Let $p > 2$ be a prime, and let $\mathfrak{P} = \mathbb{Z}[\zeta_m](p)$; this is a prime ideal of $\mathbb{Z}[\zeta_m]$ lying over p in the sense that $\mathfrak{P} \cap \mathbb{Z} = (p)$. Now, if only the representation could be realized over $\mathbb{Z}[\zeta_m]$, we could simply reduce mod \mathfrak{P} and obtain a representation over an extension of \mathbb{F}_p . However, this is either not always true or just not known to be true by the authors. To fix the problem, one first localizes at \mathfrak{P} ; that is, we consider the ring $\mathbb{Z}[\zeta_m]_{\mathfrak{P}}$ of all fractions with the property that the denominator lies outside \mathfrak{P} ; this is a principal ideal domain with a single prime (and maximal) ideal \mathfrak{P} . In this case, the representation can be realized over $\mathbb{Z}[\zeta_m]_{\mathfrak{P}}$, as this PID generates the whole field as its field of fractions (see [CR06, §73.6]). Now we can reduce mod \mathfrak{P} ; the result is a matrix realization over the field $\mathbb{Z}[\zeta_m]_{\mathfrak{P}}/\mathfrak{P}$; it is easy to check that this is an extension of the field $\mathbb{F}_p = \mathbb{Z}/(p)$. Furthermore, the dimension of this extension field is the multiplicative order of p modulo m (the same as the extension of the splitting field of the polynomial $X^m - 1$ over \mathbb{F}_p). This immediately gives rise to a representation over the field \mathbb{F}_q with $q = p^{\text{ord}_m p} \leq p^{\phi(m)}$. We remark that this process preserves irreducibility, and induces a complete decomposition of $\mathbb{F}_p[G]$ into irreducible representations.